

REMARKS

Claims 1-17 and 19-27 are pending in the present application. Claim 18 is canceled. Reconsideration of the claims is respectfully requested.

Amendments are made to the specification to correct errors and to clarify the specification. No new matter is added by any of the amendments to the specification.

I. 35 U.S.C. § 112, Second Paragraph

The Office Action rejects claims 18-27 under 35 U.S.C. § 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter, which applicants regard as the invention. This rejection is respectfully traversed.

The Office Action states:

In claim 18, and therefore its dependent claims, it is not entirely clear how the first switch connects the first memory to the other elements (see Figure 1, and note the amendment changing "memory controller" to --memory--).

Claim 18 is canceled. Claims 19, 22, and 23 are amended to include the limitations previously presented in claim 18 and to overcome the objection. Claims 19, 22, and 23 are amended to recite a first memory controller and that the switch connects the host adapter, the processor, the first memory controller, and the drive adapter, as previously presented in original claim 18. Therefore, Applicants respectfully request withdrawal of the rejection of claims 19-27 under 35 U.S.C. § 112, second paragraph.

II. 35 U.S.C. § 102, Anticipation

The Office Action rejects claims 1-18 under 35 U.S.C. § 102 as being anticipated by *Chong, Jr.* This rejection is respectfully traversed.

Chong, Jr. teaches a storage architecture that provides scalable performance through independent control and data transfer paths. Commands and status information are passed to a control module in a storage controller through a control transfer path. Data may be passed from a host computer to a storage device and from the storage device to the host computer directly through a data path. The storage controller includes a

switch for routing control and data based on the messaging scheme. See *Chong, Jr.*, Abstract; col. 2, line 56, to col. 4, line 5.

Chong, Jr. also teaches a fault-tolerant configuration with scalable performance storage architecture. This architecture includes switches with additional ports to allow two hosts to simultaneously access each switch for pertinent data transfer operations involving a storage device. Each switch may send data to multiple places at the same time, such as mirrored cache memories. The described scalable performance storage architecture is shown in **Figure 5**, which is reproduced below:




Fig. 5

This architecture appears to be a configuration of separate parts, rather than discrete storage controllers. In other words, switch 1 221 is not in or a part of a first storage controller and switch 2 222 is not part of a second storage controller. Rather, the configuration of elements, when taken together, appears to form the architecture.

In contradistinction, the present invention provides a storage controller that allocates a memory buffer for a data block from a memory pool that includes a first memory on the first storage controller and a second memory on a second storage controller and allows the first storage controller to access cached data blocks stored on the second storage controller. Independent claim 1 recites:

1. A method for managing a read request, comprising:

receiving, at a first storage controller, a read request for a data block;
allocating a memory buffer for the data block from a memory pool that includes a first memory on the first storage controller and a second memory on a second storage controller, wherein the memory buffer resides in the second memory;
retrieving the data block from a storage device; and
caching the data block in the memory buffer.

Independent claim 7 recites:

7. A method for managing a read request, comprising:
receiving, at a first storage controller, a read request for a data block; and
retrieving the data block from a memory pool that includes a first memory on the first storage controller and a second memory on a second storage controller, wherein the data block resides in the second memory.

While *Chong, Jr.* does teach an architecture in which a host may access one of two of cache memories through one or two switches, *Chong, Jr.* does not teach or suggest a storage controller that is capable of allocating memory on a separate storage controller, caching a data block on a separate storage controller, and/or retrieving a cached data block stored on a separate storage controller, as recited in claims 1 and 7. That is, in the embodiments in *Chong, Jr.* where separate storage controllers are shown, there is no switch-to-switch path. See *Chong, Jr.*, Figs. 3A-3E and 4A-4B. In the embodiment of *Chong, Jr.* where multiple switches are shown with a switch-to-switch path, there are no separate and distinct controllers.

The Final Office Action alleges that the claims do not recite a first controller allocating memory on a separate storage controller. Applicants respectfully disagree. Claim 1, for example, recites a first storage controller and a second storage controller. In order for there to be a first storage controller and a second storage controller, there must be at least two distinct storage controllers. The applied reference does not make such a designation. *Chong, Jr.* actually teaches a collection of parts that operate to provide redundancy for fault tolerance. However, *Chong, Jr.* does not teach or suggest that one subset of these parts make up a first storage controller and that another subset of these parts make up a distinct second storage controller. The Office Action proposes drawing a dashed box around components to illustrate the breadth of the claim language.

Applicants contend that dashed lines cannot be arbitrarily drawn around components of an overall system to meet the limitations of claims.

The applied reference does not teach or suggest each and every claim limitation; therefore, claims 1 and 7 are not anticipated by *Chong, Jr.* Since claims 2-5 and 8-10 depend from claims 1 and 7, the same distinctions between *Chong, Jr.* and the invention recited in claims 1 and 7 apply for these claims. Additionally, claims 2-5 and 8-10 recite other additional combinations of features not suggested by the reference. Claims 18-22 recite subject matter addressed above with respect to claims 1-5 and 7-10 and are allowable for the same reasons.

More particularly, as to claims 2, 3, 8, 9, 12, and 13, the Office Action states:

With respect to claims 2, 8 and 12, *Chong, Jr.* discloses that the first controller includes a first switch (221 in Fig. 5) and the second controller includes a second switch (222 in Fig. 5).

With respect to claims 3, 9 and 13, *Chong, Jr.* discloses that the first switch and the second switch are coupled using a switch-to-switch path [note the bus or path between the switches in Fig. 5, for example].

Office Action, dated December 1, 2003. Applicant respectfully disagrees. While *Chong, Jr.* does teach a first and second switch and a switch-to-switch path, *Chong, Jr.* does not teach or fairly suggest that a first switch in a first storage controller and a second switch in a second storage controller wherein the switches of the first storage controller and the second storage controller are coupled using a switch-to-switch path. The applied reference does not teach or suggest each and every claim limitation; therefore, claims 2, 3, 8, 9, 12, and 13 are not anticipated by *Chong, Jr.*

Still more particularly, claim 4 recites receiving a read request at a first storage controller and storing the data block in a second memory at a second storage controller via a switch-to-switch path. As stated above, *Chong, Jr.* does not teach or suggest allocating a memory buffer for the data block from a memory pool that includes a first memory on the first storage controller and a second memory on a second storage controller, wherein the memory buffer resides in the second memory, as recited in claim 1. While *Chong, Jr.* teaches a collection of parts that operate to provide redundancy for fault tolerance, *Chong, Jr.* does not teach or suggest that one subset of these parts make up a first storage controller and that another subset of these parts make

up a distinct second storage controller. Thus, it follows that *Chong, Jr.* does not teach the further limitation of receiving a read request for a data block at a first storage controller and storing the data block in the memory of a second storage controller, as recited in claim 4.

Claim 10 recites retrieving the data block from the second memory at the second controller using the switch-to-switch path. As sated above, *Chong, Jr.* does not teach or suggest retrieving the data block from a memory pool that includes a first memory on the first storage controller and a second memory on a second storage controller, wherein the data block resides in the second memory. While *Chong, Jr.* teaches a collection of parts that operate to provide redundancy for fault tolerance, *Chong, Jr.* does not teach or suggest that one subset of these parts make up a first storage controller and that another subset of these parts make up a distinct second storage controller. Thus, it follows that *Chong, Jr.* does not teach the further limitation of receiving a read request for a data block at a first storage controller and reading the data block from the memory of a second storage controller, as recited in claim 10.

As to claim 11, the present invention provides a storage controller that allocates a memory buffer for a data block from a memory pool that includes a first memory on the first storage controller and a second memory on a second storage controller and allows the first storage controller to access cached data blocks stored on the second storage controller. Independent claim 11 recites:

11. A method for managing a write request, comprising:
 - receiving, at a first storage controller, a write request for a data block;
 - allocating a primary data buffer for the data block in a first memory and a mirror data buffer for the data block in a second memory, wherein the first memory resides on one of the first storage controller and a second storage controller and the second memory resides on the other of the first storage controller and the second storage controller;
 - storing write data for the data block in the primary data buffer; and
 - mirroring the write data in the mirror data buffer.

While *Chong, Jr.* does teach an architecture in which a host may mirror data blocks in two cache memories through one or two switches, *Chong, Jr.* does not teach or suggest a storage controller that is capable of allocating memory on a separate storage controller, and mirroring a cached data block between the storage controller and a separate storage

controller, as recited in claim 11. That is, in the embodiments in *Chong, Jr.* where separate storage controllers are shown, there is no switch-to-switch path. See *Chong, Jr.*, Figs. 3A-3E and 4A-4B. In the embodiment of *Chong, Jr.* where multiple switches are shown with a switch-to-switch path, there are no separate and distinct controllers. See *Chong, Jr.*, Figure 5.

The applied reference does not teach or suggest each and every claim limitation; therefore, *Chong, Jr.* does not anticipate claim 11. Since claims 12-17 depend from claim 11, the same distinctions between *Chong, Jr.* and the invention recited in claim 11 apply for these claims. Additionally, claims 12-17 recite other additional combinations of features not suggested by the reference.

More particularly, claim 15 recites the second memory resides on the second storage controller and the step of mirroring the write data in the mirror data buffer comprises storing the write data in the mirror data buffer via the switch-to-switch path. As sated above, *Chong, Jr.* does not teach or suggest allocating a primary data buffer for the data block in a first memory and a mirror data buffer for the data block in a second memory, wherein the first memory resides on one of the first storage controller and a second storage controller and the second memory resides on the other of the first storage controller and the second storage controller, as recited in claim 11. While *Chong, Jr.* teaches a collection of parts that operate to provide redundancy for fault tolerance, *Chong, Jr.* does not teach or suggest that one subset of these parts make up a first storage controller and that another subset of these parts make up a distinct second storage controller. Thus, it follows that *Chong, Jr.* does not teach the further limitation of receiving a write request for a data block at a first storage controller and mirroring the write data in the mirror data buffer on a second storage controller, as recited in claim 15.

III. Allowable Subject Matter

Applicants thank the Examiner for the indication of allowable subject matter. Claims 19, 22, and 23 are amended to be in independent form and to overcome the rejection under 35 U.S.C. 112, second paragraph, as recommended by the Examiner. Consequently, claims 19-27 are believed to be allowable.

IV. Conclusion

It is respectfully urged that the subject application is patentable over the prior art of record and is now in condition for allowance.

The Examiner is invited to call the undersigned at the below-listed telephone number if in the opinion of the Examiner such a telephone conference would expedite or aid the prosecution and examination of this application.

DATE: July 20, 2004

Respectfully submitted,



Stephen R. Tkacs
Reg. No. 46,430
Yee & Associates, P.C.
P.O. Box 802333
Dallas, TX 75380
(972) 367-2001
Agent for Applicants